

A Tool for Scheduling THEMIS Observations

Gregg Rabideau*, Steve Chien*, David McLaren*, Russell Knight*,
Saadat Anwar**, Greg Mehall**, Philip Christensen**

* Jet Propulsion Laboratory, California Institute of Technology, USA
e-mail: firstname.lastname@jpl.nasa.gov

**School of Earth and Space Exploration, Arizona State University, USA
e-mail: {sanwar,mehall, phil.christensen}@asu.edu

Abstract

We describe a scheduling system intended to assist in the development of instrument data acquisitions for the THEMIS instrument, onboard the Mars Odyssey spacecraft. This tool creates observations of both (a) targeted geographical regions of interest and (b) general mapping observations, while respecting spacecraft constraints such as data volume, observation timing, visibility, lighting, season, and science priorities. This tool therefore must address both geometric and state/timing/resource constraints. We describe a tool that maps geometric polygon overlap constraints to set covering constraints using a grid-based approach. These set covering constraints are then incorporated into a greedy optimization scheduling algorithm incorporating operations constraints to generate feasible schedules. The resultant tool generates schedules of hundreds of observations per week out of potential thousands of observations. This tool is currently under evaluation by the THEMIS observation planning team at Arizona State University.

1 Introduction

In April of 2001, NASA launched the Mars Odyssey spacecraft carrying several instruments including the Thermal Emission Imaging System (THEMIS) for the purpose of collecting multi-spectral data of the surface of Mars. Since the start of science mapping in February of 2002, THEMIS has provided a vast dataset that is used in a wide range of scientific studies.

With this success, however, comes the complex task of selecting science targets for the instrument as the Martian surface quickly passes underneath. In January 2010, the planets aligned in such a way to allow THEMIS to collect data at a higher rate than previously achieved. While increasing the success of the mission, this also compounded the problem of selecting observations from the many viewing opportunities. We address this problem using automated planning and scheduling technology that efficiently selects THEMIS observations which satisfy the complex set of requirements from the spacecraft, instrument, and

scientists. In this paper, we describe our automated process and results, and compare them with the current process and results. Developed in collaboration by the Artificial Intelligence Group of the Jet Propulsion Laboratory and the THEMIS science planning team at the Arizona State University, the THEMIS Observation Scheduling Tool (TOST) is currently being evaluated by the science planning team at ASU.

Specifically, within TOST, we divide the problem into three primary steps: swath generation, campaign generation, and target selection.

1. In the first step, a ground track of the spacecraft is used to compute the regions of the Martian surface viewable by the THEMIS instrument at each point in time. These regions are represented as time-tagged polygons and the scheduling problem can be viewed as selecting a subset of the potential observation polygons to maximize a prioritized score of science coverage goals while respecting spacecraft operations constraints.

2. In the second step, campaigns are generated to represent the prioritized imaging requests of the scientists. In some cases, “targeted observations”, a region-of-interest (ROI) is identified on the surface, along with specific observational parameters (lighting, season, etc.). However, there is also a general science goal of constructing a global of the surface of Mars under a range of conditions (e.g., a global map at 2pm local time, global map during spring). These are so called “mapping observations.” Because the campaign areas may not be contiguous, the campaign goals are represented as operations on polygons including intersection, union, and negation. Each of these (potentially non-contiguous) regions also has a priority and the type of data requested (an instrument mode constraint).

3. Finally, the last step is to select observations from #1 above that maximize a priority weighted score defined by the science campaigns in #2 above. In this selection, relevant spacecraft operations constraints must be met such as: data volume, instrument on-time, observation separation, command storage, and others. For this, we use an adaptation of the Compressed Large-scale Activity Scheduler Planner (CLASP) [Knight and Chien 2006] that uses squeaky wheel optimization (SWO) [Fox 1996, Joslin &

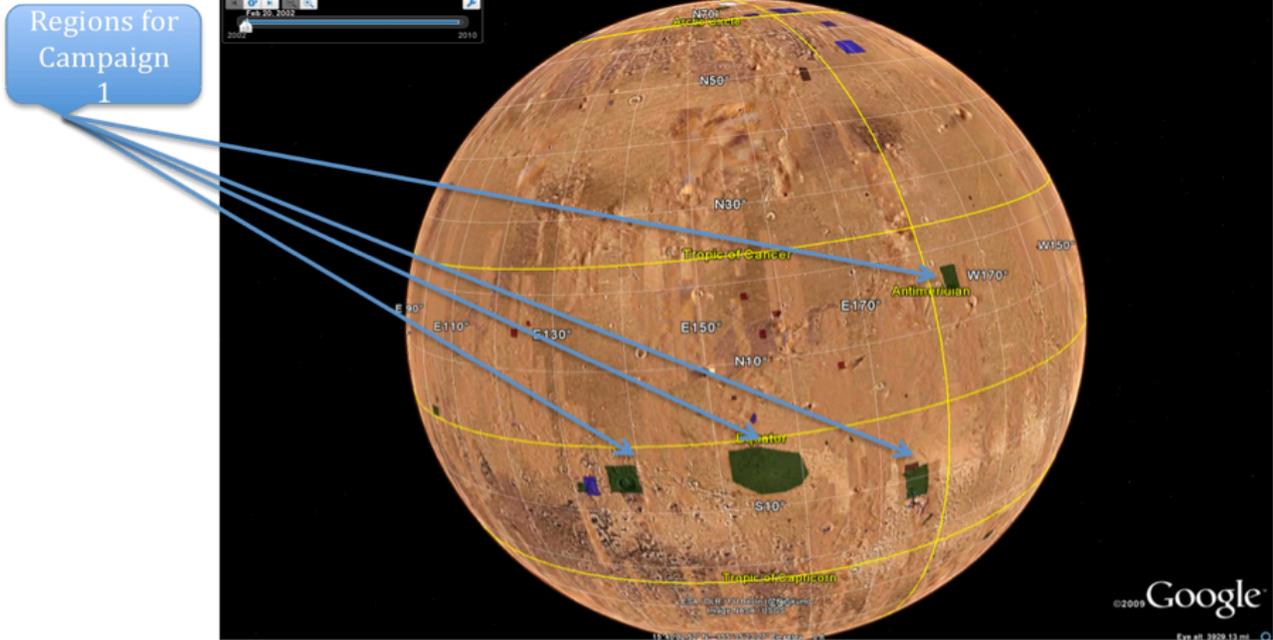


Figure 1: Mineral Search Regions on Mars

Clements 1999] iterative heuristic approach to select observations.

In the remainder of this paper we describe the problem formulation, scheduling algorithm, and project status.

2 Swath Generation

In swath generation, we first retrieve the Mars ground track of the Odyssey spacecraft by querying a server running at Arizona State University (ASU) that uses Navigation Ancillary Information Facility (NAIF) [NAIF] orbital data and the SPICE toolkit to calculate coordinates for a given time range. Next, polygons are created from ground track points representing the area on the surface of Mars that is viewable by the instrument. THEMIS has two observation modes - infrared (IR) and visible (VIS). IR and VIS have different swaths,

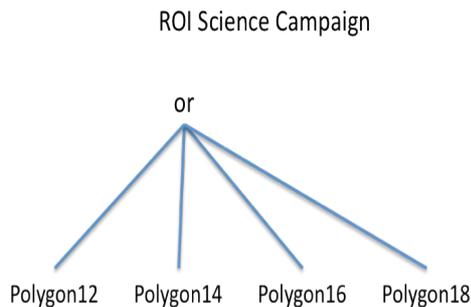


Figure 2: Tree representation of ROI Science Campaign

operations constraints, operations modes, and data rates. Consequently, a separate swath is generated for each instrument and mode. For example, VIS has a swath width of 18.4 km and IR has a swath width of 32.0 km. The IR instrument can operate in several modes, acquiring up to 10 spectral bands of data where more bands of data means that the instrument has a higher data rate. The VIS instrument can acquire up to 5 spectral bands and typically is capturing less than 5 bands due to data volume restrictions. The VIS instrument can also acquire data at 18, 36, and 72 meters per pixel resolution.

Additionally, certain instrument-mode combinations are not desired. For example, acquiring VIS images during the night would not generate useful data. Therefore, all night segments are removed from VIS swaths. Certain other overflight-specific viewing constraints are also important to the scientists. These constraints include: day versus night, restrictions on season of year (also called Ls or day of year), and local time (Lt). Because these constraints depend on the time of the observation we construct additional special instrument swaths for these potential observations.

These instrument coverage polygons and their time tags are combined to make the instrument swaths that are passed as input to observation selection process (see below).

3 THEMIS Campaigns

The THEMIS science team uses the construct of campaigns and regions of interest (ROI) to represent the desire to acquire imagery of regions of the Martian

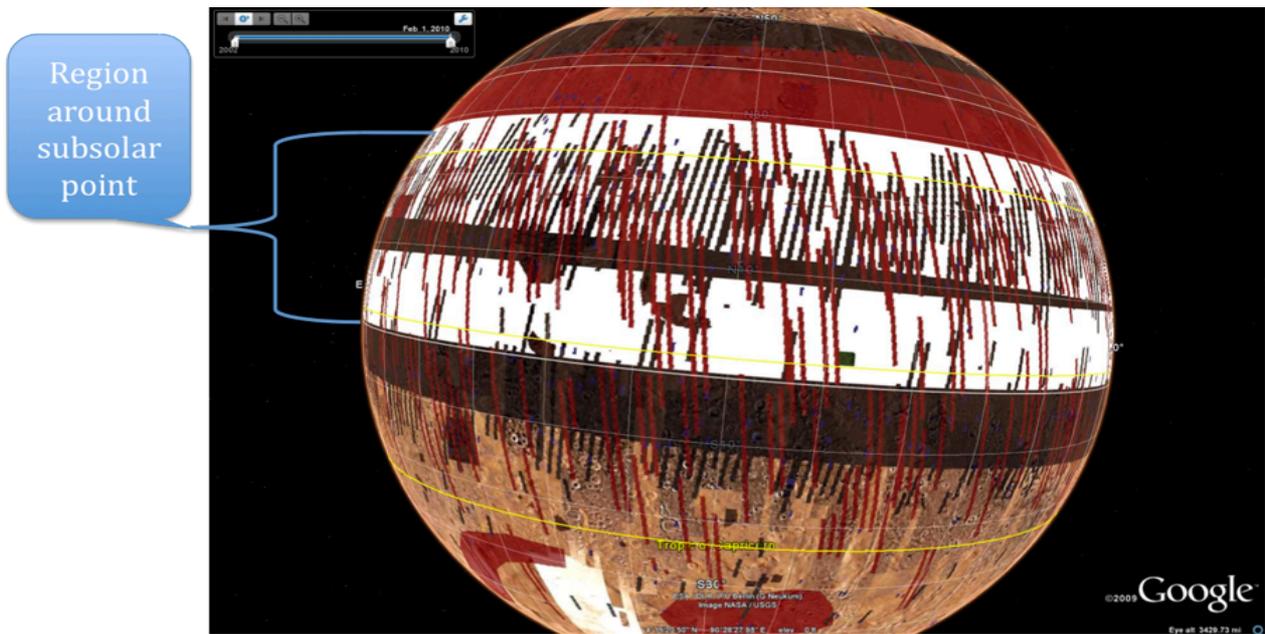


Figure 3: Mapping Campaign around Subsolar Point

surface. Within the TOST In campaign generation, we use three types of campaigns identified by the scientists: ROI, mapping, and repeat campaigns.

1. A “targeted” or ROI campaign represents a request to map a small area of the Martian surface under prescribed conditions. In an ROI campaign, the scientist specifies a polygon on the surface of Mars, along with the instrument mode to be used and optional constraints on when data can be acquired (e.g., seasonal, local time, relative

position of the sun). When an ROI has timing constraints, the ROI uses a special swath that contains only those segments that fall within the required time range. Otherwise, the ROI uses the general swath for the requested instrument mode.

2. Mapping campaigns represent the science goal of mapping the entire Martian surface under prescribed conditions (such as 3pm Local Time, or within 20 degrees of the subsolar point). As such, mapping represents a sustained campaign to map vast areas of the Martian surface with the goal of leaving no uncovered areas. In mapping campaigns, we start by constructing polygons for all of the previously acquired observations that meet the mapping campaign constraints. These areas are excluded from the requested mapping area. Also, because new (planned but not yet executed or recently acquired) observations may not have been evaluated for data quality, these observations are excluded (e.g. provisionally presumed good quality).
3. Finally, for repeat campaigns, observations are requested for areas that were previously acquired with the same set of request parameters. Repeat campaigns are treated similarly to ROI campaigns except that they are requested to be imaged every overflight that meets the side constraints.

Mapping Campaign

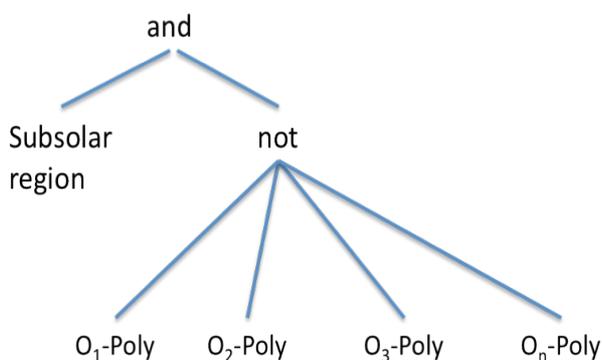


Figure 4: Tree representation of Mapping Campaign Example

All campaigns are assigned priorities based on preferences specified by the scientists. For example, mapping campaigns are assigned priorities partly based on how close previous observations have met an

assigned target allocation for the specified data type. In the end, the generated campaigns and priorities are passed as inputs to CLASP.

Because science campaigns often represent non-contiguous regions of the Martian surface, science campaigns require a more complex representation than polygons. Campaigns are represented as decision trees with internal nodes representing and/or/negation combinations and leaves representing spatial constraints (e.g. latitude north of 10 degrees north). With this semantics a subtree represents a (possibly non contiguous) region on the surface of Mars. There are currently 29 active campaigns represented in TOST.

For example, to represent a campaign to search for a mineral might involve acquiring images over several non-contiguous areas on the surface of Mars. These might be represented as the polygons shown in Figure 1 and as the campaign tree shown below in Figure 2.

As another example, a mapping campaign might wish to map the areas with the best solar illumination (as represented by the subsolar point where the sun is strongest on the surface of Mars). If some of those areas have already been acquired through prior observations (as indicated in Figure 3), the regions might be represented as shown in Figures 3 & 4 below.

4 Observation Selection

In order to assess areal coverage, CLASP uses a gridded representation of regions. In this representation, the planetary surface is represented by a set of roughly equidistant grid points with separation D . Specifically, grid points would exist along lines of latitude that are spaced distance D apart. Along these lines there would be grid points spaced D apart, surrounding the globe.

This gridded representation allows CLASP to compute overlap between regions very efficiently. With this representation rather than computing polygon overlap on a surface directly the computation simply intersection in grid point sets. Gridded overlap computation is bit set intersection and is $O(n)$ theoretically where n is the number of points in the grid but in practice these bit vector operations are in practice effectively constant time. Polygon overlap computation is $O(n \log n)$ theoretically and in practice $O(n)$ where n is the number of points defining the polygons.

For the TOST application, we use a 3200 gridpoints around the Mars Equator which converts to ~6.63 km between grid points and 3.29M grid points to represent the Martian globe.

CLASP-TOST currently considers a total of 17 instrument modes. Note that some of these instrument modes subsume others (e.g. IR observation with Band 1 is subsumed by IR observation with Bands 1 & 2). In these cases TOST must consider that one

observation may satisfy multiple requests simultaneously.

CLASP first computes the intersection points between instrument swaths and campaigns. This is done by iterating through instrument swath points and for applicable points that appear in one or more ROI's, creating a "potential observation" record for each such point, for each such ROI, if it requires a unique instrument mode. For example, if campaign1 requires 10 band IR for a point and campaign2 requires 4 band VIS, then two observation records are created. If both campaign3 and campaign4 require 10 band IR only one observation record is created. Each observation record is then accorded the highest priority from each of its campaigns.

The observation selection problem is the following:

Given

a set of potential observation records

$O = \{o_1 \dots o_n\}$

a set of regions of interest $R = \{r_1 \dots r_m\}$

a set of instrument swaths $I = \{i_1 \dots i_n\}$

Where $\forall o_i \in O \exists (r_i, i_i) \text{ grid}(o_i)$

$\in \text{grid}(r_i) \wedge \text{grid}(o_i) \in \text{grid}(i_i)$

a scoring function $U(r_i) \rightarrow \text{real}$

a constraint function $C(S) \rightarrow T, F$

where $S \subseteq O$ and C is True if

S satisfies spacecraft constraints

Select a set of observations A

To maximize $\sum a \in A U(a)$

subject to $C(A) \rightarrow T$

CLASP/TOST currently validates a number of operations constraints:

Observation spacing – with the exception of VIS images embedded within IR images, after one THEMIS observation has completed, THEMIS observations must be spaced with a minimum temporal separation. This can be represented as a simple temporal distance constraints between observations.

Observation length – because THEMIS IR observations are based on calibration made at the beginning of the observation, THEMIS IR observations that are too long result in poor quality science data near the end of the observation. Therefore IR observations are limited in length (time duration). This can be represented as a temporal distance constraint between the start and end of any THEMIS IR observations.

Onboard Storage – due to limited storage onboard the Odyssey spacecraft, the amount of data taken by THEMIS is limited by this storage capacity until renewed as indicated by a provided downlink schedule.

Command buffer – there is also a limitation on the number of command slots for uploaded sequences onboard the spacecraft. THEMIS must not exceed this limit at any time - restricting number of observations between command uplinks (command uplinks are effectively exogenous events).

CLASP uses squeaky wheel optimization, an iterative heuristic approach to optimization. In this approach, a simple greedy selection (scheduling) method is used iteratively with tweaks to the inputs to this algorithm made each iteration.

For the TOST application, each iteration is a call to SWO_inner below and consists of iterating through the potential observation records in order of decreasing priority. If the instrument swath can be added without violating any spacecraft operations constraints it is added. Else the observation record is discarded and the next observation record is considered.

Whenever an observation record is added to the schedule, CLASP must compute which additional observation records are also implied to be in the schedule (the Propagate function below). This propagation occurs based on two checks. The instrument swath polygon associated with the selected observation record may include multiple grid points. For any of these grid points (and the original selected grid point) any observation record whose instrument mode is subsumed by the selected instrument mode is also covered. For example, if the selected observation record specified instrument mode “VIS 36m resolution 4 band” it subsumes the request for “VIS 72m resolution 3 band.” An instrument mode I1 subsumes another instrument mode I2 if I1 contains all of the bands contains in I2 and I1 is at the same or higher resolution than I2. This subsumption is implemented by a lookup table.

The result of SWO_inner is a set of observation records A such that C(A) is satisfied.

For the TOST application the outer loop of SWO consists of first initializing the observation record priorities to the priority of the parent science campaigns. Then SWO_outer repeatedly calls SWO_inner to produce a set of selected observation records A. As long as a progress metric is satisfied, we increment the priority of all observation records that did not make it into the current schedule A, and re-run. This proceeds a number of iterations and the best schedule (scored by initial priorities) is returned.

SWO outer loop

Initialize priorities of all observation records to the priorities of their parent science campaigns

While progress made

SWO_inner → A

For each o in O – A increment the priority of o

repeat

SWO_inner

O = all candidate observation records

B = {}

For each o in O in decreasing priority order

If C(B+o+Propagate(o)) = True

B := B + o + Propagate (o)

5 Algorithm Performance – Theoretical and Practical for THEMIS

The theoretical algorithmic performance of CLASP-TOST is as follows [Knight 2005a, Knight & Smith 2005b]:

Swath generation to compute grid points for the instrument swath:

$O(gP)+P'$

where g is the number of grid points in the bounding box containing the polygon

and P is the number of grid points defining the polygon

and P' is the number of grid points not in the bounding box containing the polygon

Campaign creation:

$O((GlogG)T)$

where G is the number of grid points in the universal area (in our case the Mars grid) and GlogG represents the cost of performing the operation to merge to leaves grid points based on an “and” or “or” and

T is the number of internal nodes in the tree, and represents the number of times we have to perform the merge operation.

Each run of the SWO_inner algorithm requires $O(I)$ observation insertions and calls to C.

For TOST, observation insertion is $O(N)$ where N is the number of timeline events in the schedule.

The number of SWO_inner calls I is user specifiable and is a small number (e.g. 10).

Practically speaking, THEMIS science planners work on two schedules per week, each of 3-4 Earth days at a time. However this is in part due to the challenge of manually considering so many observations and operations constraints. Because of the automation, there is interest in constructing scheduled of 7 days for analysis purposes and CLASP-TOST has been tested on a one week planning horizon.

Each day translates into hundreds of thousands of map grid points that must be evaluated. On-board storage for science data is the primary factor limiting the THEMIS observation volume, allowing only a few hours of observation time each day. Below we show the runtime performance of CLASP-TOST on a four and seven day schedules (64 bit Red Hat Linux, dual core 2.4GHz AMD Opteron, 16GB RAM, -O2 optimization).

Schedule Duration	4 Earth Days	7 Earth Days
# of observations in generated schedule	421	758
Time to generate instrument Swaths	~2 minutes	~3 minutes
Time to generate science ROI's	~2 minutes	~2 minutes
CLASP Initialization	~10 minutes	~20 minutes
CLASP # of iterations, time per iteration	3 x ~15.5 minutes	3 x ~83 minutes
Total CLASP Time	56 minutes	4 hours and 31 minutes

6 THEMIS Operations and Evaluation Status

The TOST system is designed for use by science planners to generate an initial set of observations. The science planners will then evaluate and manually edit the schedule with any changes desired. These schedules are then used to generate THEMIS command sequences which are uploaded and execute onboard the Odyssey Spacecraft.

Operationally, CLASP-TOST schedules are generated as KML files. These KML files can be loaded by Google Earth[Google Earth] or JMARS [JMARS] for visual inspection. From JMARS the files can be saved out as selected observations for later command generation.

7 Discussion, Related Work, Conclusions

Spacecraft operations have been a major area of application for automated planning and scheduling. Numerous space missions have used automated planning & scheduling on the ground to enable significant operational efficiencies including the Hubble Space Telescope [Johnston et al. 1993], space shuttle refurbishment [Deale et al. 1994], shuttle payload operations [Chien et al. 1999], The Modified Antarctic Mapping Mission [Smith et al. 2002], Mars Exploration Rovers [Bresina et al. 2005], Earth Observing One (EO-1) [Chien et al. 2005a] Mars Express [Cesta et al. 2005], and Orbital Express [Chouinard et al. 2008]. Automated planning has even flown as a technology demonstration on the Deep Space One (DS1) Mission

[Muscettola et al. 1998] and as the primary operations system on 3CS [Chien et al. 2001] and EO-1 [Chien et al. 2005b]. However, all of the above applications focused on the state, resource, and timing aspects of mission operations rather automating both the spatial coverage as well the state and resource reasoning. A notable exception is [Knight and Hsu 2009] which also uses the CLASP system.

This work represents a preliminary implementation of a scheduling system designed to assist in the scheduling of spatial campaign observations for the THEMIS instrument of the Mars Odyssey Mission. Future work includes both tool enhancements and scheduling algorithm analysis. In the tool enhancement area we would like to investigate means of explaining why observations are or are not selected. This could include information on the science campaigns that motivated selection of an observation or computation of which selected observations are in conflict with a proposed observation. Of course, further evaluation by the ASU THEMIS science planning team is a top priority.

Further analysis of the THEMIS scheduling problem and the TOST-CLASP tool would also be useful. Evaluation against other search algorithms besides squeaky wheel would be useful. In particular, approaches that could guarantee optimal solutions such as branch and bound would be ideal. However the large problems sizes for THEMIS scheduling may pose a problem for near exhaustive search. Further analysis of the key pre-processing and scheduling complexities are needed. Derivation of upper bounds on optimal schedules via solution of relaxed versions of the problem (as in [Chien et al. 2010]) seems to offer some promise for analysis. Better characterization of the problem sizes for the THEMIS scheduling would also be helpful.

8 Conclusions

This paper has described a mission planning tool for the THEMIS instrument currently flying onboard the Mars Odyssey Spacecraft. This tool, called TOST, is used to generate candidate observation schedules. TOST first constructs spatial observation candidates for both the THEMIS instrument and THEMIS science regions of interest. These geometric constraints are then combined with spacecraft operations constraints by the CLASP planner using the Squeaky Wheel heuristic greedy optimization algorithm. The CLASP-TOST tool is currently under evaluation by the THEMIS science planning team.

References

J. L. Bresina, A. K. Jonsson, P. H. Morris, K. Rajan, Activity planning for the mars exploration rovers, Proceedings 15th International Conference on

Automated Planning & Scheduling, Monterey, CA, 2005.

S. Chien, G. Rabideau, J. Willis, T. Mann, "Automating Planning and Scheduling of Shuttle Payload Operations," *Artificial Intelligence Journal* 114 (1999) 239-255.

S. Chien, B. Cichy, A. Davies, D. Tran, G. Rabideau, R. Castano, R. Sherwood, D. Mandl, S. Frye, S. Shulman, J. Jones, S. Grosvenor, "An Autonomous Earth Observing Sensorweb," *IEEE Intelligent Systems*, May-Jun 2005, pp. 16-24.

S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, D. Mandl, S. Frye, B. Trout, S. Shulman, D. Boyer, "Using Autonomy Flight Software to Improve Science Return on Earth Observing One, *Journal Aerospace Computing, Information, & Communication*, April 2005, AIAA.

S. Chien, D. Tran, G. Rabideau, S. Schaffer, D. Mandl, S. Frye, "Timeline-based space operations scheduling with external constraints," *Proceedings International Conference on Automated Planning & Scheduling*, Toronto, Canada, 2010.

C. Chouinard, R. Knight, G. Jones, D. Tran, D. Koblick, *Automated & Adaptive Mission Planning for Orbital Express*, Space Operations, Heidelberg, Germany, 2008.

Christensen, P. R., Jakosky, B. M., Kieffer, H. H., Malin, M. C., McSween, Jr., H. Y., Nealon, K. Mehall, G. L., Silverman, S. H., Ferry, S., and Caplinger, M., 'The Thermal Emission Imaging System (THEMIS) for the Mars 2001 Odyssey Mission', *Space Science Reviews*, 110, 85–130, 2004.

M. Deale, M. Yvanovich, D. Schnitzius, D. Kautz, M. Carpenter, M. Zweben, G. Davis, and B. Daun, "The Space Shuttle Ground Processing Scheduling System," in *Intelligent Scheduling*, Zweben & Fox (eds.), Morgan Kaufman, San Francisco, 1994.

B. Fox, An Algorithm for scheduling improvement by schedule shifting, Technical Report 96.5.1. McDonnell Douglas Aerospace, Houston, 1996. See also Scheduling Optimizer B. R. Fox US Patent 5,890,134 (1999).

Google Earth, <http://earth.google.com/index.html>

JMARS Java Mission Planning and Analysis for Remote Sensing, <http://jmars.asu.edu/>

M. D. Johnston, R. Henry, A. Gerb, M. Giuliano, B. Ross, N. Sanidas, S. Wissler and J. Mainar, "Improving the Observing Efficiency of Hubble Space Telescope," *AIAA Aerospace*, AIAA Press, 1993.

D. Joslin, and D. Clements, "Squeaky wheel optimization," *Journal of Artificial Intelligence Research* 10:353-373, 1999.

R. Knight, "Solving the Constrained Coverage Problem using Flow Networks as Linear Program Approximations," PhD. Dissertation, University of California, Los Angeles, 2005.

R. Knight, B. Smith, "Optimally Solving Nadir Observation Scheduling Problems," *International Symposium on Artificial Intelligence, Robotics, and*

Automation in Space, Munich, Germany, September 2005.

R. Knight, S. Chien. "Producing Large Observation Campaigns Using Compressed Problem Representations," *Proceedings of the Fifth International Workshop on Planning and Scheduling for Space*, Baltimore, MD, October 2006.

R. Knight, S. Hu. "Compressed Large-scale Activity Scheduling and Planning (CLASP) applied to DESDynI," *Proceedings of the Sixth International Workshop in Planning and Scheduling for Space*, Pasadena, CA, July 2009.

Mars Odyssey Mission Home Page, <http://marsprogram.jpl.nasa.gov/odyssey/>

N. Muscettola, P. Nayak, B. Pell, B. C. Williams, "Remote Agent: to boldly go where no AI system has gone before," *Artificial Intelligence* v103, # 1-2, pp. 5–47, 1998.

Navigation Ancillary Information Facility, NAIF/SPICE, <http://naif.jpl.nasa.gov/naif/index.html>

B.D. Smith, B.E. Engelhardt, and D.H. Mutz, "The RADARSAT-MAMM Automated Mission Planner," *AI Magazine*, v. 23, # 2, 2002. pp. 25–36.